arXiv:1512.05974v2 [cs.DS] 13 May 2016

Improved Balanced Flow Computation Using Parametric Flow

Omar Darwish[a,*], Kurt Mehlhorn[a]

[a]*Max-Planck-Institute für Informatik, Saarbrücken, Germany*

Abstract

We present a new algorithm for computing balanced flows in equality networks arising in market equilibrium computations. The current best time bound for computing balanced flows in such networks requires $O(n)$ maxflow computations, where n is the number of nodes in the network [Devanur et al. 2008]. Our algorithm requires only a single parametric flow computation. The best algorithm for computing parametric flows [Gallo et al. 1989] is only by a logarithmic factor slower than the best algorithms for computing maxflows. Hence, the running time of the algorithms in [Devanur et al. 2008] and [Duan and Mehlhorn 2015] for computing market equilibria in linear Fisher and Arrow-Debreu markets improve by almost a factor of n.

Keywords: Algorithms, Network Flow, Market Equilibrium, Balanced Flow, Parametric Flow

1. Introduction

Balanced flows are an important ingredient in equilibrium computations for the linear Fisher and Arrow-Debreu markets. In the Arrow-Debreu market model, we have n agents, where each agent owns some goods and a utility function over the existing set of goods. The goal is to assign prices to all the goods such that the market clears, meaning that all the goods are sold and each agent spends all of his budget on only goods with maximum utility. These prices are called the equilibrium prices and the state of the market with such prices is the equilibrium state we want to reach. Under some assumptions, it was shown in [7] that such prices exist. The Fisher model is a simplification of the Arrow-Debreu model, where each agent comes with a budget instead of the endowment of goods [8]. The linear Arrow-Debreu market is a special case where the utility functions are linear, and the same applies to the linear Fisher model. Later on, we will refer to agents as buyers.

The problem of finding the equilibrium prices in these market models has been studied for a long time, see [9] and [10] for surveys. The concept of a balanced flow was introduced in the first combinatorial algorithm [1] for finding the equilibrium prices in the linear Fisher market. Later, balanced flows were used in [2] and [11] for combinatorial algorithms for computing equilibrium prices in the linear Arrow-Debreu market.

The equilibrium algorithms mentioned above compute balanced flows in a special type of flow network, known as *equality network*. An *equality network* is a bipartite flow network where we have a set of buyers B adjacent to the source node s and a set of goods C adjacent to the sink node t. Hence, the vertex set of the network is $\{s,t\} \cup B \cup C$. Also, we assume that the number of buyers is equal to

*Corresponding author

Email addresses: odarwish@mpi-inf.mpg.de (Omar Darwish), mehlhorn@mpi-inf.mpg.de (Kurt Mehlhorn)

the number of goods. Each buyer b_i has a positive budget e_i and each good c_j has a positive price p_j. The edge set is defined as follows:

1. An edge (s, b_i) with capacity e_i for each $b_i \in B$.
2. An edge (c_i, t) with capacity p_i for each $c_i \in C$.
3. All edges running between B and C have infinite capacity. Also, we assume that each buyer is connected to at least one good and similarly each good is connected to at least one buyer.

A flow in this network corresponds to the money flow from buyers to goods. So, an amount of x units flowing from buyer b_i to good c_j indicates that b_i spends x units on good c_j. We refrain from relating the equality network to the market equilibrium computation as this is not important for understanding our algorithm.

Balanced flows are defined for equality networks in [1]. We denote the set of buyers by $B = \{b_1, b_2, ..., b_n\}$. The surplus $r(b_i)$ of the buyer b_i with respect to a maximum flow f is the residual capacity of the edge (s, b_i) (capacity of this edge minus the flow running through it). The surplus vector is the vector of the surpluses of all buyers $(r(b_1), r(b_2), ..., r(b_n))$. The balanced flow is the maximum flow that minimizes the two-norm of the surplus vector among all maximum flows of the network. In [1], it is shown that a balanced flow can be computed with $O(n)$ maxflow computations. This is the best time bound achieved up till now for computing balanced flows.

The name "balanced flow" comes from the fact that it is the maximum flow in the equality network that balances the surpluses as much as possible. Additionally, it is shown in [1] that the surplus vector is unique for all balanced flows in a given equality network. From this fact and other properties shown in [1] and [2], we can deduce the following characterization of balanced flows in equality networks:

1. Buyers can be partitioned into maximal disjoint blocks based on their surpluses, $B_1, B_2, ..., B_h$ having surpluses $r_1 > r_2 > ... > r_h \geq 0$. Buyers in block B_i have the surplus r_i. The uniqueness of this partition follows from the uniqueness of the surplus vector of the buyers.
2. We can assume without loss of generality that $r_h = 0$.
3. Define C_i to be the set of goods to which money flows from B_i. For $i < h$, all goods in C_i are completely sold, i.e., the edges from C_i to t are saturated. There is no flow from B_i to C_j for $j < i$. Also, there are no edges from B_i to C_j for $j > i$.

We show how to compute balanced flows in equality networks by a single parametric flow computation, which improves the running time of the algorithms in [1] and [2] for computing market equilibria in linear Fisher and Arrow-Debreu markets by almost a factor of n. Our algorithm adds to the applications of parametric flows mentioned in [3].

In Section 2, we give an introduction to parametric network flows. In Section 3, we state our algorithm for computing balanced flows, analyse its running time, and prove its correctness.

2. Parametric Network Flows

The parametric network flow problem [3] is a generalization of the standard network flow problem, in which the arc capacities are not fixed, but are functions of a real valued single parameter λ. More specifically, we consider the following parametric problem:

1. The arc (s, b_i) has capacity $\max(0, e_i - \lambda)$.

2. The capacity of all other arcs is constant.

3. The parameter λ decreases from a large value down to zero.

We define the *min-cut capacity function* $\kappa(\lambda)$ as the capacity of a minimum cut in the network as a function of the parameter λ [3]. Among the cuts of capacity $\kappa(\lambda)$, let $(X(\lambda), \overline{X(\lambda)})$ be the cut with the smallest sink side[1] $\overline{X(\lambda)}$. It is shown in [3] that $X(\lambda) \subseteq X(\lambda')$ for $\lambda \geq \lambda'$.

This implies that we have at most $n-1$ distinct minimum cuts. Under the assumption of linearity of the arc capacity functions of λ, $\kappa(\lambda)$ is a piecewise-linear concave function with at most $n-2$ breakpoints. A breakpoint is a value of λ where the slope of $\kappa(\lambda)$ changes.

As long as the minimum cut is the same, decreasing λ will only result in increasing the flow linearly in the current segment of $\kappa(\lambda)$. [3] gives an algorithm for computing all breakpoints of $\kappa(\lambda)$ which runs in $O(nm\log(n^2/m))$ worst-case time complexity, for a network with n nodes and m edges; they call it the *breakpoint algorithm*. Additionally, they state that this algorithm can be easily augmented to store for each vertex $v \notin \{s,t\}$ the breakpoint at which v moves from the sink side to the source side of a minimum cut of minimum sink size without altering the time bound. We will refer to the breakpoint algorithm with the stated augmentation as the *augmented breakpoint algorithm*. We will use this *augmented breakpoint algorithm* in our algorithm for computing balanced flows.

The algorithm in [3] extends the standard preflow algorithm by Goldberg and Tarjan [12] for solving the maximum flow problem. Its running time is only by the factor $\log(n^2/m)$ slower than the running time of the best maxflow algorithms [4, 5, 6].

3. Improved Algorithm for Balanced Flow Computation

In this section, we describe our algorithm for computing a balanced flow in a given equality network N and prove its correctness. The intuition behind our algorithm is as follows. For $\lambda = \infty$, all edges out of s have capacity zero. Since all edges into t have positive capacity, the minimum cut will have all buyers on the sink side. Then, step by step, at λ values corresponding to $r_1, ..., r_h$, the blocks $B_1, ..., B_h$ of buyers will move from the sink side of the minimum cut to the source side. The behavior of the function $\kappa(\lambda)$ changes at these breakpoints, since at these surpluses some buyers will not be able to push more flow to the goods side. Hence the flow will stop at this point at these buyers, while continuing to increase for other buyers, and so on. An example for the evolution of the minimum cut is shown in Figure 1. We will next give the details.

3.1. The Algorithm

Our algorithm works on an equality network N as follows:

1. Construct the parametric network PN from N. First, construct PN as a copy of N. Next, change the capacity of each edge (s, b_i) in PN from e_i to $\max(0, e_i - \lambda)$.

2. Run the *augmented breakpoint algorithm* on PN, in order to find all the breakpoints of $\kappa(\lambda)$ for PN and record for each buyer b_i the breakpoint at which b_i moves from the sink side to the source side of a minimum cut of minimum sink size; let it be λ_i for each b_i. Note that, initially all buyers are on the sink side of the minimum cut as shown in Figure 1. This will be explained in details in the analysis.

[1] With respect to set inclusion, there is a unique minimum cut of capacity $\kappa(\lambda)$ with smallest sink side.

3. Construct the network N' which is a copy of N, but with changing the capacity of each edge (s,b_i) from e_i to $e_i - \lambda_i$. Note that, we are setting the capacity of each edge running from s to buyer b_i to be e_i minus the parameter value of the breakpoint where this buyer moves in the minimum cut from the sink side to the source side. This breakpoint of buyer b_i is the surplus of this buyer in the balanced flow of N, as we will show in the analysis.

4. Compute the maximum flow in N'. It is a balanced flow of the initial network N.

3.2. Complexity and Correctness Analysis

First, we analyze the worst-case time complexity of our algorithm. We can easily see that steps (1) and (3) run in time complexity equivalent to the size of the network, so steps (2) and (4) are dominating, in terms of time complexity. As explained in the previous section, step (2) runs in $O(nm\log(n^2/m))$ worst-case time complexity, for an equality network with n nodes and m edges. For step (4), it needs one maxflow computation, so it is dominated by step (2). Hence, the overall worst-case time complexity of our algorithm is $O(nm\log(n^2/m))$.

Now, we prove the correctness of the algorithm. Recall from the previously stated characterization of balanced flows that we have a unique partition of buyers $B_1, B_2, ..., B_h$ having surpluses $r_1 > r_2 > ... > r_h = 0$, where buyers in block B_i have the surplus r_i. Also, recall that we define C_i to be the set of goods to which money flows from B_i. Using this characterization, we can deduce the following lemmas.

Lemma 3.1. *For $i < h$ and $\lambda \in (r_{i+1}, r_i]$, the minimum cut of minimum sink size corresponding to $\kappa(\lambda)$ in the parametric network PN is $(s \cup (B_1 \cup B_2 \cup ... \cup B_i) \cup (C_1 \cup C_2 \cup ... \cup C_i), t \cup (B_{i+1} \cup B_{i+2} \cup ... \cup B_h) \cup (C_{i+1} \cup C_{i+2} \cup ... \cup C_h))$.*

Proof. Let $(X, \overline{X})$ be a minimum cut of minimum sink size of capacity $\kappa(\lambda)$ and let f be a balanced flow in N.

We will first prove that the set of buyers $D = B_1 \cup B_2 \cup ... \cup B_i \subseteq X$. Notice that D consists of all buyers whose surplus with respect to f is greater than or equal to λ. Assume for the sake of a contradiction that $K = D \cap \overline{X}$ is non-empty. Let L be the set of goods in $\overline{X}$ that are connected by an edge to a buyer in K; buyers in K may also be connected to goods in $C \cap X$ and f may carry nonzero flow from K to $C \cap X$. From the characterization of balanced flows, f does not carry flow from $B \setminus D$ to L. Since $(X, \overline{X})$ is a cut, there can be no edges from buyers in $D \setminus K$ to L; note that such edges have infinite capacity. We conclude that in f all flow to the goods in L comes from the buyers in K. Since all goods in L are completely sold with respect to f (note that $i < h$), we conclude

$$\sum_{c_j \in L} p_j \leq \sum_{b_j \in K} (e_j - r_j) \leq \sum_{b_j \in K} (e_j - \lambda),$$

where the last inequality follows from $\lambda \leq r_i$. Thus moving $K \cup L$ from $\overline{X}$ to X will not increase the capacity of the cut and will reduce the sink side of the cut. We have now shown that $D \subseteq X$. Since the edges from D to $\cup_{j \leq i} C_j$ have infinite capacity and each good in $\cup_{j \leq i} C_j$ is incident to one such edge, we also have $\cup_{j \leq i} C_j \subseteq X$.

We will next prove that $\overline{D} = B_{i+1} \cup B_{i+2} \cup ... \cup B_h = B \setminus D \subseteq \overline{X}$. Assume for the sake of a contradiction, that $F = X \cap \overline{D}$ is nonempty. Let M be the goods that are connected to F. Since the edges from F to M have infinite capacity, we have $M \subseteq X$, and hence the edges running from M to t contribute to the capacity of the cut. We know that the capacity of these edges is greater than or equal

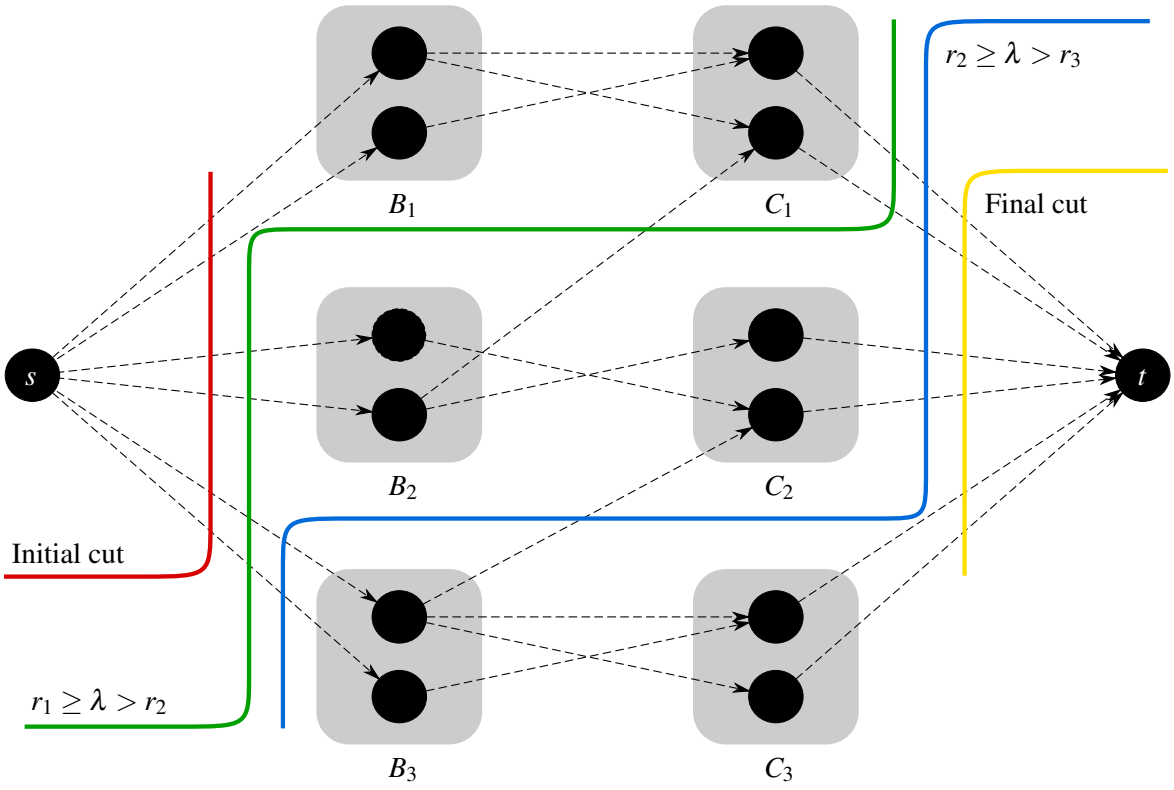

Figure 1: An example of an equality network in the form of a balanced flow partition. We have 3 blocks B_1, B_2, and B_3 of buyers having surpluses $r_1 > r_2 > r_3$, where buyers in block B_i have surplus r_i. The figure shows how the minimum cut is changing throughout our algorithm based on the parameter value λ. Initially at $\lambda = \infty$, the minimum cut includes only s in the source side. As λ decreases, the minimum cut will change based on the surpluses. At $\lambda = r_1$, B_1 and C_1 move from the sink side to the source side, then similarly B_2 and C_2 move from the sink side to the source side at $\lambda = r_2$. The final minimum cut includes only t in the sink side.

to $\sum_{b_j \in F} e_j - r_j$, as in the balanced flow the goods in M receive $\sum_{b_j \in F} e_j - r_j$ units of flow from F. Since the surplus of each of the buyers in F with respect to the balanced flow is less than λ, we have $\sum_{b_j \in F} e_j - \lambda < \sum_{b_j \in F} e_j - r_j \leq \sum_{c_j \in M} p_j$. Thus moving F and M from the source to the sink side of the cut will reduce the capacity of the cut (convince yourself that moving F and M to the sink gives a cut). This proves that all buyers in $\overline{D}$ belong to the sink side of the minimum cut. Naturally, all goods adjacent to buyers in $\overline{D}$ without edges from other buyers will also be in the sink side, hence the sink side of the minimum cut contains the nodes $(t \cup (B_{i+1} \cup B_{i+2} \cup \ldots \cup B_h) \cup (C_{i+1} \cup C_{i+2} \cup \ldots \cup C_h))$. This concludes the proof. □

Lemma 3.2. *The breakpoints of $\kappa(\lambda)$ for the network PN correspond to the surpluses of the blocks in any balanced flow of N.*

Proof. Let us decrease λ from infinity to zero. We see from lemma 3.1 that the first breakpoint we will meet is $\lambda = r_1$. At this point all the buyers in B_1 will move in the minimum cut of minimum sink size corresponding to $\kappa(\lambda)$ from the sink side to the source side. Similarly, $\lambda = r_2$ will be our second breakpoint, where all buyers from B_2 will move from the sink side to the source side of the minimum cut. The same argument will apply for all the breakpoints. □

Theorem 3.3. *The algorithm computes a balanced flow of the equality network N, having n nodes on each side and m edges, in $O(nm\log(n^2/m))$ worst-case time complexity.*

Proof. We proved the worst-case time complexity in the beginning of this section. Concerning the correctness of the algorithm, we proved in Lemma 3.2 that the breakpoints of $\kappa(\lambda)$ for the network *PN* correspond to the surpluses of the blocks in any balanced flow of *N*. Hence, computing these surpluses will tell us the surplus of each buyer in the balanced flow network and that's what we do in steps (1) and (2) of the algorithm. Now, step (3) enforces these surpluses of each buyer by subtracting this surplus from the capacity of each edge running from the source to the buyer, constructing the network N'. So, when we compute the maximum flow in N', all edges from the source node to the buyers will be saturated, hence each buyer will have the corresponding correct surplus. Also, this surplus vector is unique and corresponds to the balanced flow, hence this maximum flow in N' is a balanced flow of *N*. □

Acknowledgements

We would like to thank the anonymous reviewer for his/her constructive comments, in particular, with respect to Figure 1.

References

[1] N. R. Devanur, C. H. Papadimitriou, A. Saberi, V. V. Vazirani, Market Equilibrium via a Primal–dual Algorithm for a Convex Program, J. ACM 55 (5) (2008) 22:1–22:18.

[2] R. Duan, K. Mehlhorn, A Combinatorial Polynomial Algorithm for the Linear Arrow-Debreu Market, Information and Computation 243 (2015) 112–132, a preliminary version appeared in ICALP 2013, LNCS 7965, pages 425-436.

[3] G. Gallo, M. D. Grigoriadis, R. E. Tarjan, A Fast Parametric Maximum Flow Algorithm and Applications, SIAM J. Comput. 18 (1989) 30–55.

[4] J. Cheriyan, T. Hagerup, A Randomized Maximum-Flow Algorithm, SIAM Journal of Computing 24 (2) (1995) 203–226.

[5] V. King, S. Rao, R. Tarjan, A Faster Deterministic Maximum Flow Algorithm, Journal of Algorithms 17 (3) (1994) 447 – 474, ISSN 0196-6774.

[6] J. B. Orlin, Max Flows in O(Nm) Time, or Better, in: Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing, STOC '13, ACM, New York, NY, USA, ISBN 978-1-4503-2029-0, 765–774, 2013.

[7] K. J. Arrow, G. Debreu, Existence of an Equilibrium for a Competitive Economy, Econometrica 22 (3) (1954) 265–290, ISSN 00129682, 14680262.

[8] I. Fisher, Mathematical Investigations in the Theory of Value and Prices, Ph.D. thesis, Yale University, 1891.

[9] N. R. Devanur, J. Garg, L. A. Végh, A Rational Convex Program for Linear Arrow-Debreu Markets, CoRR abs/1307.8037.

[10] B. Codenotti, S. Pemmaraju, K. Varadarajan, The Computation of Market Equilibria, SIGACT News 35 (4) (2004) 23–37, ISSN 0163-5700.

[11] R. Duan, J. Garg, K. Mehlhorn, An Improved Combinatorial Algorithm for the Linear Arrow-Debreu Market, In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '16, pages 90–106. SIAM, 2016.

[12] A. Goldberg, R. E. Tarjan, A new Approach to the Maximum-Flow Problem, Journal of the ACM 35 (1988) 921–940.